

What is claimed is:

1. A method of performing a join, the method comprising:

identifying a join that identifies a right table and a left table, the right table including a plurality of right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions, the left table including one or more left table rows;

determining that a partitioning column of the right table is specified in an equality condition of the join;

preparing the left table for joining one or more rows of the left table with one or more rows of the right table, wherein preparing the left table comprises:

selecting one of the plurality of left table rows; and

generating a partition number for the selected one of the plurality of left table rows;

for a left table row for which a partition number was generated:

identifying a right table partition with a partition number that matches the generated partition number; and

joining, using a product join, the one or more right table rows in the matching partition with the selected one of the left table rows if one or more join conditions are satisfied.

2. The method of Claim 1, wherein preparing the left table for joining further comprises generating a spool of the left table.

3. The method of Claim 1, wherein preparing the left table for joining further comprises:

generating a partition number for each of the plurality of left table rows; and

sorting the left table rows by partition number.

4. The method of Claim 3, wherein preparing the left table further comprises:

placing the left table rows associated with the same partition number into a cache.

5. The method of Claim 4, further comprising:

joining, using a product join, the one or more right table rows in the matching partition with the each of the left table rows in the cache if one or more join conditions are satisfied.

6. The method of Claim 1, wherein preparing the left table for joining further comprises duplicating the left table to each of the processing modules.

7. A method of performing a join, the method comprising:

identifying a join that identifies a left table and a right table, the left table including a plurality
5 of left table rows, each of the plurality of left table rows grouped into one of a plurality
of left table partitions, the right table including one or more right table rows, each of the
plurality of right table rows grouped into one of a plurality of right table partitions;

determining that the left table and the right table are joined on equality constraints;

determining a relationship between a partitioning expression associated with the left table and a
10 partitioning expression associated with the right table;

determining that each left table partition matches at least one right table partition based on the
relationship between the partitioning expressions associated with the left table and right
table; and

joining the one or more left table rows of each matching left table partition with the one or
15 more right table rows of the at least one matching right table partition if one or more
join conditions are satisfied.

8. The method of Claim 7, wherein determining that the left table and the right table are joined on
equality constraints comprises:

determining that the join specifies an equality constraint between each primary index column of
20 the left table and a corresponding primary index column of the right table; and

determining that the join specifies an equality constraint between each partitioning column of
the left table and a corresponding partitioning column of the right table.

9. The method of Claim 7, wherein determining the relationship between the partitioning expressions
comprises:

25 determining that a mapping of the plurality of left table partitions to the plurality of right table
partitions is a one to one relationship.

10. The method of Claim 9, wherein determining that at least one left table partition matches at least
one right table partition comprises:

matching each partition of the left table with only one partition of the right table.

11. The method of Claim 9, wherein joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least one matching right table partition comprises performing a merge join.

12. The method of Claim 7, wherein determining a relationship between the partitioning expressions
5 comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship.

13. The method of Claim 12, wherein determining that at least one left table partition matches at least one right table partition comprises:

10 matching each partition of the left table with at least two partitions of the right table.

14. The method of Claim 13, wherein joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions comprises:

15 joining the one or more left table rows of each partition of the left table with the one or more rows of the at least two matching partitions of the right table.

15. The method of claim 7, wherein determining a relationship between the partitioning expressions comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a many-to-many relationship.

20 16. The method of claim 15, wherein determining that at least one left table partition matches at least one right table partition comprises:

matching each partition of the left table with at least two partitions of the right table; and
matching each partition of the right table with at least two partitions of the left table.

17. A method of performing a join, the method comprising:

identifying a join that identifies a left table and a right table, the left table including a plurality of left table rows, each of the plurality of left table rows grouped into one of a plurality of left table partitions, the right table including one or more right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions;

determining that the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table;

determining that inequality conditions exist between each partitioning column of the left table and a corresponding partitioning column of the right table;

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship;

determining that at least one left table partition matches at least two right table partitions based on the relationship between the partitioning expressions associated with the left table and right table; and

joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions if one or more join conditions are satisfied.

18. The method of claim 17, wherein determining that inequality conditions exist between each partitioning column of the left table and a corresponding partitioning column of the right table

comprises:

determining that the inequality conditions comprise one or more inequality expressions and one or more expressions that are not inequality expressions.

19. A method of performing a join, the method comprising:

identifying a join that identifies a first table and a second table, the first table including a plurality of first table rows, each of the plurality of first table rows grouped into one of a plurality of first table partitions, the second table including one or more second table rows;

determining that a partitioning column of the first table is specified in an equality condition of the join;

preparing the second table for joining one or more rows of the second table with one or more rows of the first table, wherein preparing the second table comprises:

selecting one of the plurality of second table rows; and

generating a partition number for the selected one of the plurality of second table rows;

for a second table row for which a partition number was generated:

identifying a first table partition with a partition number that matches the generated partition number;

joining, using a product join, the one or more first table rows in the matching partition with the selected one of the second table rows if one or more join conditions are satisfied.

20. A computer program, stored on a tangible storage medium, for use in performing a join, the program including executable instructions that cause a computer to:

identify a join that identifies a right table and a left table, the right table including a plurality of right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions, the left table including one or more left table rows;

determine that a partitioning column of the right table is specified in an equality condition of the join;

prepare the left table for joining one or more rows of the left table with one or more rows of the right table, wherein preparing the left table comprises:

select one of the plurality of left table rows; and

generate a partition number for the selected one of the plurality of left table rows;

for a left table row for which a partition number was generated:

identify a right table partition with a partition number that matches the generated partition number; and

join, using a product join, the one or more right table rows in the matching partition with the selected one of the left table rows if one or more join conditions are satisfied.

21. The computer program of Claim 20, wherein, when preparing the left table for joining, the computer generates a spool of the left table.

22. The computer program of Claim 20, wherein, when preparing the left table for joining, the computer:

generates a partition number for each of the plurality of left table rows; and

sorts the left table rows by partition number.

23. The computer program of Claim 20, wherein, when preparing the left table, the computer:

places the left table rows associated with the same partition number into a cache.

24. The computer program of Claim 23, further comprising executable instructions that cause the computer to:

join, using a product join, the one or more right table rows in the matching partition with the each of the left table rows in the cache if one or more join conditions are satisfied.

25. The computer program of Claim 20, wherein, when preparing the left table for joining, the computer duplicates the left table to each of the processing modules.

26. A computer program, stored on a tangible storage medium, for use in performing a join, the program including executable instructions that cause a computer to:

- 5 identify a join that identifies a left table and a right table, the left table including a plurality of left table rows, each of the plurality of left table rows grouped into one of a plurality of left table partitions, the right table including one or more right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions;
- 10 determine that the left table and the right table are joined on equality constraints;
- 10 determine a relationship between a partitioning expression associated with the left table and a partitioning expression associated with the right table;
- 10 determine that each left table partition matches at least one right table partition based on the relationship between the partitioning expressions associated with the left table and right table; and
- 15 join the one or more left table rows of each matching left table partition with the one or more right table rows of the at least one matching right table partition if one or more join conditions are satisfied.

27. The computer program of Claim 26, wherein, when determining that the left table and the right table are joined on equality constraints, the computer:

- 20 determines that the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table; and
- 20 determines that the join specifies an equality constraint between each partitioning column of the left table and a corresponding partitioning column of the right table.

28. The computer program of Claim 26, wherein, when determining the relationship between the partitioning expressions, the computer:

- 25 determines that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to one relationship.

29. The computer program of Claim 28, wherein, when determining that at least one left table partition matches at least one right table partition, the computer:

- 30 matches each partition of the left table with only one partition of the right table.

30. The computer program of Claim 28, wherein, when joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least one matching right table partition, the computer performs a merge join.

31. The computer program of Claim 26, wherein determining a relationship between the partitioning expressions comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship.

32. The computer program of Claim 31, wherein, when determining that at least one left table partition matches at least one right table partition, the computer:

matches each partition of the left table with at least two partitions of the right table.

33. The computer program of Claim 32, wherein, when joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions, the computer:

joins the one or more left table rows of each partition of the left table with the one or more rows of the at least two matching partitions of the right table.

34. The computer program of Claim 26, wherein, when determining a relationship between the partitioning expressions, the computer:

determines that a mapping of the plurality of left table partitions to the plurality of right table partitions is a many-to-many relationship.

35. The computer program of Claim 34, wherein, when determining that at least one left table partition matches at least one right table partition, the computer:

matches each partition of the left table with at least two partitions of the right table; and
matches each partition of the right table with at least two partitions of the left table.

36. A computer program, stored on a tangible storage medium, for use in performing a join, the program including executable instructions that cause a computer to:

5 identify a join that identifies a left table and a right table, the left table including a plurality of left table rows, each of the plurality of left table rows grouped into one of a plurality of left table partitions, the right table including one or more right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions;
determines that the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table;
determines that inequality conditions exist between each partitioning column of the left table
10 and a corresponding partitioning column of the right table;
determines that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship;
determines that at least one left table partition matches at least two right table partitions based on the relationship between the partitioning expressions associated with the left table
15 and right table; and
joins the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions if one or more join conditions are satisfied.

37. The computer program of claim 36, wherein, when determining that inequality conditions exist
20 between each partitioning column of the left table and a corresponding partitioning column of the right table, the computer:

determines that the inequality conditions comprise one or more inequality expressions and one or more expressions that are not inequality expressions.

38. A method of performing a join, the method comprising:

identifying a join that identifies a second table and a first table, the second table including a plurality of second table rows, each of the plurality of second table rows grouped into one of a plurality of second table partitions, the first table including one or more first table rows, each of the plurality of first table rows grouped into one of a plurality of first table partitions;

determining that the second table and the first table are joined on equality constraints;

determining a relationship between a partitioning expression associated with the second table and a partitioning expression associated with the first table;

determining that each second table partition matches at least one first table partition based on the relationship between the partitioning expressions associated with the second table and first table; and

joining the one or more second table rows of each matching second table partition with the one or more first table rows of the at least one matching first table partition if one or more join conditions are satisfied.

39. A database system including:

a massively parallel processing system including:

one or more nodes;

a plurality of CPUs, each of the one or more nodes providing access to one or more
5 CPUs;

a plurality of data storage facilities each of the one or more CPUs providing access to
one or more data storage facilities;

a process for execution on the massively parallel processing system for performing a join, the
process including:

10 identifying a join that identifies a right table and a left table, the right table including a
plurality of right table rows, each of the plurality of right table rows grouped
into one of a plurality of right table partitions, the left table including one or
more left table rows;

determining that a partitioning column of the right table is specified in an equality
15 condition of the join;

preparing the left table for joining one or more rows of the left table with one or more
rows of the right table, wherein preparing the left table comprises:

selecting one of the plurality of left table rows; and

generating a partition number for the selected one of the plurality of left table rows;

20 for a left table row for which a partition number was generated:

identifying a right table partition with a partition number that matches the
generated partition number; and

joining, using a product join, the one or more right table rows in the matching
partition with the selected one of the left table rows if one or more join
25 conditions are satisfied.

40. The database system of Claim 39, wherein preparing the left table for joining further comprises
generating a spool of the left table.

41. The database system of Claim 39, wherein preparing the left table for joining further comprises:

generating a partition number for each of the plurality of left table rows; and

30 sorting the left table rows by partition number.

42. The database system of Claim 41, wherein preparing the left table further comprises:
placing the left table rows associated with the same partition number into a cache.

43. The database system of Claim 42, wherein the process further comprises:
joining, using a product join, the one or more right table rows in the matching partition with the
5 each of the left table rows in the cache if one or more join conditions are satisfied.

44. The database system of Claim 39, wherein preparing the left table for joining further comprises
duplicating the left table to each of the processing modules.

45. A database system including:

a massively parallel processing system including:

10 one or more nodes;
a plurality of CPUs, each of the one or more nodes providing access to one or more
CPUs;
a plurality of data storage facilities each of the one or more CPUs providing access to
one or more data storage facilities;

15 a process for execution on the massively parallel processing system for performing a join, the
process including:

identifying a join that identifies a left table and a right table, the left table including a
plurality of left table rows, each of the plurality of left table rows grouped into
one of a plurality of left table partitions, the right table including one or more
20 right table rows, each of the plurality of right table rows grouped into one of a
plurality of right table partitions;

determining that the left table and the right table are joined on equality constraints;

determining a relationship between a partitioning expression associated with the left
table and a partitioning expression associated with the right table;

25 determining that each left table partition matches at least one right table partition based
on the relationship between the partitioning expressions associated with the left
table and right table; and

30 joining the one or more left table rows of each matching left table partition with the one
or more right table rows of the at least one matching right table partition if one
or more join conditions are satisfied.

46. The database system of Claim 45, wherein determining that the left table and the right table are joined on equality constraints comprises:

determining that the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table; and

determining that the join specifies an equality constraint between each partitioning column of the left table and a corresponding partitioning column of the right table.

47. The database system of Claim 45, wherein determining the relationship between the partitioning expressions comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to one relationship.

48. The database system of Claim 47, wherein determining that at least one left table partition matches at least one right table partition comprises:

matching each partition of the left table with only one partition of the right table.

49. The database system of Claim 47, wherein joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least one matching right table partition comprises performing a merge join.

50. The database system of Claim 45, wherein determining a relationship between the partitioning expressions comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship.

51. The database system of Claim 50, wherein determining that at least one left table partition matches at least one right table partition comprises:

matching each partition of the left table with at least two partitions of the right table.

52. The database system of Claim 51, wherein joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions comprises:

joining the one or more left table rows of each partition of the left table with the one or more rows of the at least two matching partitions of the right table.

53. The database system of claim 45, wherein determining a relationship between the partitioning expressions comprises:

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a many-to-many relationship.

5 54. The database system of claim 53, wherein determining that at least one left table partition matches at least one right table partition comprises:

matching each partition of the left table with at least two partitions of the right table; and
matching each partition of the right table with at least two partitions of the left table.

55. A database system including:

a massively parallel processing system including:

one or more nodes;

a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;

a plurality of data storage facilities each of the one or more CPUs providing access to one or more data storage facilities;

a process for execution on the massively parallel processing system for performing a join, the process including:

identifying a join that identifies a left table and a right table, the left table including a plurality of left table rows, each of the plurality of left table rows grouped into one of a plurality of left table partitions, the right table including one or more right table rows, each of the plurality of right table rows grouped into one of a plurality of right table partitions;

determining that the join specifies an equality constraint between each primary index column of the left table and a corresponding primary index column of the right table;

determining that inequality conditions exist between each partitioning column of the left table and a corresponding partitioning column of the right table;

determining that a mapping of the plurality of left table partitions to the plurality of right table partitions is a one to many relationship;

determining that at least one left table partition matches at least two right table partitions based on the relationship between the partitioning expressions associated with the left table and right table; and

joining the one or more left table rows of the at least one matching left table partition with the one or more right table rows of the at least two matching right table partitions if one or more join conditions are satisfied.

56. The database system of claim 55, wherein determining that inequality conditions exist between each partitioning column of the left table and a corresponding partitioning column of the right table comprises:

determining that the inequality conditions comprise one or more inequality expressions and one or more expressions that are not inequality expressions.

57. A method of performing a join, the method comprising:

identifying a join that identifies a second table and a first table, the second table including a plurality of second table rows, each of the plurality of second table rows grouped into one of a plurality of second table partitions, the first table including one or more first table rows, each of the plurality of first table rows grouped into one of a plurality of first table partitions;

determining that the join specifies an equality constraint between each primary index column of the second table and a corresponding primary index column of the first table;

determining that inequality conditions exist between each partitioning column of the second table and a corresponding partitioning column of the first table;

determining that a mapping of the plurality of second table partitions to the plurality of first table partitions is a one to many relationship;

determining that at least one second table partition matches at least two first table partitions based on the relationship between the partitioning expressions associated with the second table and first table; and

joining the one or more second table rows of the at least one matching second table partition with the one or more first table rows of the at least two matching first table partitions if one or more join conditions are satisfied.